

Daniel Ellsberg
4 Carver St.
Cambridge, Mass.

Critique of the Theory of Two-person Zero-sum Games

How well did von Neumann and Morgenstern succeed at the particular task they set themselves: "to find the mathematically complete principles which define 'rational behavior' for the participants in a social economy, and to derive from them the general characteristics of that behavior."? (p. 31). At the time of the second edition of *The Theory of Games*, the authors were satisfied that with respect to the special "two-person zero-sum" game (though not for more general cases) they had obtained "a precise theory... which gives complete answers to all questions."

Curiously, this claim has received little careful attention; it has been ignored, by unsympathetic listeners, or uncritically accepted by expositors and game theorists. Even the one unfavorable critic who has published on this subject, Carl Kaysen, has accepted this particular conclusion within the limits of the authors' assumptions, though he has gone on to question those assumptions. For the rest of the published articles, Arrow has expressed the common view: "The theory of rational behavior in zero-sum two-person games can therefore be regarded as definitely solved."

The question is important for two reasons. The solution to the two-person game is made the essential foundation of the theory of more general games, including "oligopoly" games (which are generally non-zero-sum); in fact, nearly every individual theorem in the general theory relies on this initial solution. Second, apart from its role in game-theory, any such results as the authors have claimed would constitute a solution to an important case of rational choice under uncertainty. It would thus introduce the concept of rationality into many economic situations involving choice under uncertainty, where it has been previously undefined and orthodox theory has been correspondingly "indeterminate."

only the VNM theory
two different words
what do you mean by them

Do they say that this
sentence is the nec.
and sufficient condition

What does
this mean?

as it happens if

$$n = 1$$

the V.N. $< n$ definition is the
old defⁿ of "nationality"

2.

This paper will be most concerned with the relationship of their results to the general problem of defining rational choice under uncertainty. Von Neumann and Morgenstern have indeed produced a meaningful definition (i.e., one which can be applied unambiguously), and we need not minimize that achievement, but this in itself is clearly not enough. In order to evaluate their conclusions, it will be necessary to establish some criteria for a "satisfactory" definition. The authors offer several suggestions toward this. For example, they state that they expect a solution to consist of "a complete set of rules of behavior in all conceivable situations." (p. 33). Moreover, "all conceivable situations" must be interpreted to include "those where 'the others' behaved irrationally, in the sense of the standards which the theory will set for them." (p. 32). That is, "the rules of rational behavior must provide definitely for the possibility of irrational conduct on the part of others."

These properties are obviously not sufficient conditions for a satisfactory definition; they would apply equally, for example, to definitions of irrational or imitative behavior. We approach closer to sufficiency if we require that the new definition bear a close relationship to the older concept of "rationality under certainty." This condition would seem to be inescapable. Otherwise, the existence of two competing definitions of "rationality" (even though operating under different conditions) would lead to intolerable confusion. Von Neumann and Morgenstern clearly ~~xx~~ accept this restriction when they require that "rational" behavior must be in some sense more "advantageous" to a player than any other behavior, no matter whether or not his opponents followed the prescribed pattern of rationality; the "superiority" of rational behavior over any other is to be established.

1 do not believe this - you
appear to be confusing

3 problems

- 1) Ignorance
- 2) Probability - Expected value
of an outcome

i.e.: a mixed prospect

- 3) Game 04 - vs another
will or
decision unit

3.

The problem now arises of making these notions operational. How are we to decide when a particular decision is or is not "advantageous"? In the case when each available action is associated with a single, certain consequence--i.e., when an individual acts under certainty--this question has been decided. If the individual can rank the consequences in order of preference, and if these preferences are transitive, then the rule of rational behavior under certainty requires him to choose that action whose consequence he most prefers. If an observer can discover (by questioning or observation) the actions available to the actor, the consequences linked with them, and the actor's preferences among those consequences, the observer can classify the actor's decision unambiguously as "rational" or "non-rational."

incidentally what definition

But if the individual acts under uncertainty, which is to say, if each action is associated in his mind with a set of possible outcomes rather than with a unique consequence, this definition is meaningless, even if actions, ~~preferences~~ sets of outcomes, and preferences among individual outcomes¹ are known. Moreover, it has

1. ~~These~~ preferences could no longer be inferred from observations of actual choices, using "revealed preference" techniques. The individual no longer chooses an outcome but rather a set of outcomes, and he might well prefer a possible outcome of an action which he rejects to the actual outcome of the action chosen.

immediate
no ~~extension~~ extension to this sphere. If it is uncertain whether the outcome of one action will be better or worse than that of another, there is no obvious sense in which one can be said to be more advantageous.

again what do we mean here?

The concept of "rationality under certainty" has both normative and descriptive roles, the two being essentially related. Hypotheses based on the ~~concept~~ concept of rational behavior (certainty being assumed) are empirically fruitful because in fact most people try,

1 wonder?

There is the problem of setting
up a standard with certain
desirable properties - Morozuk
seems to be quite right in using
a "large number" i.e. we pick
and fix a norm that is fairly
good and observe that the deviations
are not too damaging or practically
or at least can be fixed by a small
modification such as a couple of parameters

4.

not merely tend, to follow rational principles. It seems important, if the same name and connotations are to be retained, that the new concept should ^{also} have a status as a normative principle, i.e., a rule of behavior which ~~a~~ a set of people under consideration agree that they "ought" (in some sense, not necessarily ethical) to follow. Moreover, with the same purpose of avoiding ambiguity, approximately the same set of people should accept this principle as those who are "rational under certainty."

Taking all these conditions into account,
/A principle may be considered a "useful" definition of rational choice under uncertainty if a large number of "otherwise reasonable" people ~~fixably~~ would reject, upon deliberation, any decision which was inconsistent with the principle.¹ This major criterion leaves

1. This proposition, which is crucial to later discussion, has been adapted from an unpublished paper by Jacob Marschak. It should be noted that the presence of the undefined modifier, "otherwise reasonable" does not introduce any circularity. There are any number of ways to define this notion independently. A particularly important way--perhaps even mandatory in this context--would have it include "those people who are rational under certainty" and no others.

a few questions unanswered. Must all "otherwise reasonable" people (however we define that; see footnote) accept the given principle; if not, how many? To put that question in another form, must we look for a unique principle with this property, or might we be satisfied with a set of "rational" (or simply "reasonable") principles such that all "otherwise reasonable" persons would follow one or another of them? Von Neumann and Morgenstern definitely set themselves the bolder task of finding a principle to have the status of "the" unique definition of rational behavior. This makes the test of their conclusions much easier. If we should decide that a large group of reasonable people would not reject, even after careful consideration, some decisions inconsistent with the particular

5.

principle the authors propose (whether or not their behavior were consistent with some other principle) then we must conclude that von Neumann and Morgenstern have failed. They would not have produced a principle which could be satisfactory or useful as "the" unique definition of rational choice under uncertainty.

— Let me add, at this point that just about everything you have said holds for all utility work — also you fall for the usual routine of talking about choice systems as though meaningful 'revealed preference' work had ever been done

The abstract model of the two-person zero-sum game can be described as follows. Player A selects a "strategy" i (this notion will be defined below) from the set of m strategies allowed him under the rules of the game. Simultaneously, in ignorance of A's choice, player B selects a strategy j , one of his n admissible strategies. Then, after the choices are revealed, A receives an amount a_{ij} and B receives an amount $-a_{ij}$ (i.e., B pays A an amount a_{ij}). These are the outcomes ("payoffs"), being money or a mathematical expectation of money; the subscripts indicate that each payoff is a function of both strategies. The rules of the game prescribe a pair of outcomes corresponding to each possible pair of strategies., and the sum of the outcomes is zero; what one player wins, the other loses.

~~The simplest actual game corresponding to this pattern is Matching Pennies. Each player has two strategies, Heads or Tails. For each pair of strategies a pair of outcomes is prescribed~~
"move"

In this model, each player makes but one ~~move~~. Thus, the analysis applies directly to such simple games as Matching Pennies, ~~courses,~~ in which each player chooses between the alternatives ~~moves,~~ Heads or Tails. To generalize the results to more complex games such as chess, the authors interpret the player's single move as the choice of a "strategy," a concept which they define: "a plan which specifies what choices he will make in every possible situation, for every possible actual information which he may possess at that moment in conformity with the pattern of information which the rules of the game provide for him in that case." (p. 79) When both players have chosen strategies in this sense, the outcome of the game is determined; complex games can be analyzed in "static" terms, as though thus, ~~the set of strategies is equivalent to the set of moves of the~~ each player the outcome were determined by a single choice on the ~~analysis of the game is simplified~~ part of each player.

The abstract model of the two-person zero-sum game can be described as follows. Player A selects a "strategy" i from a finite set of strategies I (from the set of n strategies allowed him under the rules of the game). Simultaneously, he ignores of A's choice, Player B selects a strategy j , one of his m admissible strategies. Then, after the choices are revealed, A receives an amount a_{ij} and B receives an amount $-a_{ij}$ (i.e., B pays A an amount a_{ij}). These are the amounts ("payoffs"), being money or a dollar's legal expectation of money; the subscripts indicate that each payoff is a function of both strategies. The rules of the game prescribe a pair of outcomes corresponding to each possible pair of strategies, and the sum of the outcomes is zero; what one player wins, the other loses.

non zero sum
game
requires at
least two

The strategies and payoff function can be represented graphically by a matrix:

Each row in this matrix represents one of A's possible strategies; each column, one of B's strategies. ~~Each field of the matrix contains an element~~
~~a number representing the outcome~~ The entire matrix can be designated (a_{ij}) , where a_{ij} is the element in the i th row and j th column; i.e. a_{ij} represents the outcome to player A specified by the rules for the pair of opposing strategies i and j . B's outcome in each case is simply $-a_{ij}$. When the elements are numbers and the meaning of the strategies is spelled out, this payoff matrix completely determines the essential features of a particular two-person zero-sum game.

Or, to approach it from the other direction, any actual game can be represented by an ^{two} appropriate payoff matrix. ~~Two games having the~~
~~same strategies but differing in the payoffs can be compared conven-~~
~~iently via their payoff matrices~~ By taking a game with given strategies and varying the matrix, it is possible to study the changing properties of new, related games, the effects of slight variations in rules. It is a major contribution of von Neumann and Morgenstern to have contributed a tool of analysis so suggestive and flexible.

To "divide the difficulties" of the analysis, the authors make some important simplifying assumptions. ^{Second,} First, the rules of the game are fixed, and are known and observed by both players. ~~xxxxxx~~ Each player knows with certainty: a) what strategies he is allowed; b) what strategies his opponent is permitted; c) the outcome corresponding to any pair of opposing strategies. In other words, ~~the~~ von Neumann and Morgenstern abstract from any uncertainties concerning the strategies or payoffs. This is an important limitation on the applicability of their results to real situations, but the assumption will be accepted in this paper.

^{First} ~~Second~~, the outcomes are ^{represented} ~~expressed~~ not in "utilities," cardinal or otherwise, but in money. The authors' digression on "cardinal utilities" ~~had~~ led to much misunderstanding on this point, but they have expressed themselves unequivocally:

"We shall therefore assume that the aim of all participants in the economic system, consumers as well as entrepreneurs, is money, or equivalently a single monetary commodity. This is supposed to be unrestrictedly divisible and substitutable, freely transferable and ~~is~~ identical, even in the quantitative sense, with whatever 'satisfaction' or 'utility' is desired by each participant." (p. 8). (the necessary property of transferability rules out the use of "von Neumann-Morgenstern cardinal utilities," defined by choices in risk-situations, even if such utilities could be defined.)

The model, then, expresses just those elements of uncertainty which von Neumann and Morgenstern wish to emphasize. Corresponding to each possible action (choice of strategy) there is a set of possible outcomes, rather than a single, certain outcome. The player does know, as assumed above, that the outcome of a particular action will be one of a given set, but which one is uncertain. The problem of which von Neumann and Morgenstern set up is to prescribe a unique "rational" choice among these sets of uncertain outcomes.

If the opponent's choice were known...

In certain special cases, it is possible to define a rule of behavior which would gain general acceptance. If the outcome of one strategy is as good or better than the outcome of another for every one of the opponent's possible strategies, the first ~~may~~ will be said to "dominate" the second. In terms of the matrix, if each element in one row is greater than the corresponding element in another row, the first strategy dominates the second. An extreme example of this is the case of non-overlapping sets of outcomes, in which every element in one set (row) is greater than every element in the second set. To choose a strategy which was dominated by another would seem very like "throwing away utility" with certainty. This suggests the rule that the "rational" player will never choose a dominated strategy; only undominated strategies will be "admissible."

Made this new definition of "reasonable".
It is intuitively clear (to use a treacherous phrase) that all, or nearly all, people who were rational under certainty would reject decisions inconsistent with this rule: so it passes our test of "reasonableness." *But if person prefers one instrument to another....*
after these two assumptions
In fact, we may regard any payoff matrix as exhibiting those strategies remaining after ~~all~~ those with non-overlapping and inferior sets of outcomes have been discarded (though a few dominated strategies may still slip by). However, this rule ~~does~~ rarely ~~may~~ dictate a unique choice. *another case: if opponent has one dominant strategy, then with assumption that opponent is "reasonable...." Keynes, Savage*

In the general case, with overlapping sets of outcomes, the first problem that arises in picking a unique rule of choice is that several principles appear as candidates, ranging from the reasonable to the doubtful. ~~Since rational choice is~~ ~~related to~~ maximizing, the ~~most plausible~~ of these rules consist of replacing each set of possible outcomes by a single number, derived from it according to the rule, and then picking the strategy corresponding to the greatest of these numbers. For example, one could represent each set by its

In certain special cases, it is possible to define a rule of behavior which would gain general acceptance. If the outcome of one strategy is as good or better than the outcome of another for every one of the opponent's possible strategies, the first may well be said to "dominate" the second. In terms of the matrix, if each element in one row is greater than the corresponding element in another row, the first strategy dominates the second. An extreme example of this is the case of non-overlapping sets of outcomes, in which every element in one set (row) is greater than every element in the second set. To choose a strategy which was dominated by another would mean very little "knowing away nothing" with certainty. This suggests the rule that the "rational" player will never choose a dominated strategy. If he knows that the "rational" player will be "rational", it is intuitively clear (to use a phrase from game theory) that all or nearly all people who play under certain conditions would reject decisions inconsistent with this rule; so it is our case of "rational" behavior. In fact, we may regard any payoff matrix as exhibiting those strategies remaining after all dominated strategies are discarded (though a few dominated strategies may still slip by). However, this rule does not necessarily dictate a unique choice. In the general case, with overlapping sets of outcomes, the first problem that arises in picking a unique rule of choice is that several principles appear as candidates, ranging from the maximization of the total payoff. Since rational choice is usually related to maximizing the total payoff of some rules consist of replacing each set of possible outcomes by a single number, derived from it according to the rule, and then picking the strategy corresponding to the greatest of these numbers. For example, one could transform each set of its

If he knows that the other is a "gambler" then he knows what will permit him

5.

greatest element, and choose that ~~extreme~~ strategy which offered the chance of the highest outcome of all; this might be called the policy of the "reckless optimist." (Modigliani; thesis p. 102). Or one might compute ^{the} ~~the~~ average, or, more generally, some weighted combination of the highest and lowest outcomes possible under each strategy, picking the strategy with the highest average. Probably ~~Perhaps~~ more players "otherwise reasonable" would favor the second rule ~~than~~ over the first, but in the absence of a principle commanding ^{is} ~~is~~ universal precedence (which neither of these ~~are~~ likely to do) there is no basis for calling the optimists "irrational." It is important to notice that the player with a blind eye to risk, while undeniably reckless, cannot be said to be "throwing away utility" even though he is playing against a ~~reasonable~~ reasonable opponent who is informed as to the payoff function. Although such an opponent would like to inflict as large an injury as possible, he is ^{handicapped} ~~handicapped~~ by the same uncertainties as the first player; he cannot know with certainty which strategy will punish the gambler. Hence, the opponent's action is uncertain, and the optimist has some chance, however small, of achieving his maximum.

Another possible procedure, which will not be discussed here, would be to "minimax regret." The point in mentioning this cluster of possible rules, each of which might claim the allegiance of a number of "reasonable" players, is to emphasize that a principle with claims to being a "unique" solution to the problem of rational choice under uncertainty must be more than "reasonable"; it must be so compelling as to ~~cause~~ cause "otherwise reasonable" players to foreswear all other principles, ~~no matter how reasonable the alternatives may appear.~~

great element, and choose that strategy which offers the
 chance of the highest outcome of all; this might be called the policy
 of the "reckless optimist." (Hedonistic; people of 1902). For one
 or, more generally, some weighted combination
 might compute the average of the highest and lowest outcomes possible
 under each strategy, picking the strategy with the highest average.
 (Probably)
 (Kahneman would give "otherwise reasonable" would favor the second
 rule than the first, but in the absence of a principle concerning
 is
 a universal procedure (which neither of these is likely to do) there
 is no basis for calling the optimist "irrational." It is important
 to notice that the player with a blind eye to risk, while (understandably)
 reckless, cannot be said to be "knowing, away reality" even though
 he is playing against a known probability and is informed as
 to the payoff function. (Although an opponent would like to in-
 flict as large an injury as possible, he is reckless in the sense
 that he does not know which strategy he cannot know which strategy
 which strategy will be chosen by the gambler. Hence, the opponent's
 action is uncertain, and the gambler has some chance, however small,
 of achieving his maximum.
 Another possible procedure, which will not be discussed here,
 would be to "minimax regret" (The point in maximizing this cluster
 of possible moves, each of which claims the advantage of a
 number of "possible" players, it is to approximate that a principle with
 gives to each "strategy" solution to the problem of choice

what are *General*
 conditions of uncertainty?

I could not
 disagree more strongly.

ONE WAY IN WHICH YOU
 MIGHT CONSIDER *minimax* IS:

How much SHOULD a man C,
 BE willing to pay A for his
 seat at the game?

6.

Von Neumann and Morgenstern do offer a principle for this role. They propose that the player should consider only the minimum element in each set of outcome: i.e., the worst that could happen to him if he played that strategy. He should then choose the strategy with the best minimum outcome. Since this particular outcome may be expressed, for player A, as $\text{Max}_i \text{Min}_j a_{ij}$, this is known as the ~~Max~~ "maximin" policy; the corresponding policy for player B is to choose the column with the lowest maximum outcome ($\text{Min}_j \text{Max}_i a_{ij}$), or "minimax."

Von Neumann and Morgenstern did not actually present this rule as applying to rational behavior under general conditions of uncertainty; they use it solely in the context of the game, in which the uncertainty arises from the interacting expectations of opposing wills. ^{they distinguish it from player against Nature} Their emphasis on the distinctiveness of this type of uncertainty seems somewhat questionable. The opponent's behavior is either uncertain or it is not; if it is, then the player's conditions of choice seem no different from those of the "player" whose outcome depends on the uncertain "strategy" of Nature (which may be assumed to be inscrutable but not hostile). The essence of the ~~normalize~~ game, in which both players choose strategies simultaneously, is that the opponent's choice is uncertain. It would seem that a rule applicable in this situation should ~~also~~ apply equally well in a situation involving uncertainty for other reasons; it is the fact, not the origin, of uncertainty which seems important. Be this as it may, it is a chief contention of von Neumann and Morgenstern that the plausibility of the maximin rule of choice in the game- situation is implied by the hostility of the opponent; it will be argued later that this claim is based on a misleading analogy from a special situation (the "minorant"

NO

7.

game, to be discussed) later) in which it is appropriate.

The primary advantage of the maximin principle is that it enables the player to avoid, with certainty, the worst possible outcome. In other terms, the worst that can happen to a player choosing

a maximin strategy is guaranteed to be better than (or at least as good as) the worst that could happen to him under any other strategy.

No matter what the opponent does, he cannot enforce the very lowest element in the matrix (except in the ~~xxxxxx~~ case when the minimum ~~xxxxxxxx~~ outcomes for each strategy are identical).

~~This property~~ There are surely many players for whom this property would not be decisive. It has been said by a proponent of the maximin rule that it "means, in effect, that that action should be chosen about which the best certain statement can be made." (theses p. 101)," but such assertions are quite misleading. Clearly one can say with just as much certainty which is the best ~~xxx~~ element in each row, or which strategy has the highest average outcome (without making any claims for the meaning or usefulness of the average). Where the "cautious pessimist" (in Modigliani's phrase) may want to know, "I can't make below a certain outcome, which is not the worst," the "optimist" may prefer to ~~xxx~~ ^{feel,} "I can make the best outcome on the board." The latter statement is just as certain, and may sound better.

This property may attract some players but not others. The "certainty" which the principle offers--of achieving a minimum outcome which is usually better than the worst--is purchased at a price. Along with the certainty that the worst possible outcome is the maximin ^{no matter what the opponent does,} goes the certainty that, the best possible outcome ~~xxxxx~~ will not exceed a particular sum. The very highest element in the row may not be much above the maximin element, and it may be very low compared to

game, to be discussed later) in which it is appropriate.

The primary advantage of the maximin principle is that it en-

ables the player to avoid, with certainty, the worst possible out-

come. In other words, the worst that can happen to a player choosing

a maximin strategy is guaranteed to be better than for at least as

good as the worst that could happen to him under any other strategy.

No matter what the opponent does, he cannot enforce the very lowest

element in the matrix (except in the maximin case when the minimum

maximin outcomes for each strategy are identical).

Unfortunately there are surely many players for whom this

property would not be decisive. It has been said by a proponent

of the maximin rule that it "means, in effect, that that action

should be chosen about which the best certain statement can be made."

(Thomas D. Igoe), "but such assertions are quite misleading. Clearly

one can say with just as much certainty which is the best and element

in each row, or which strategy has the highest average outcome.

(without making any claim for the meaning or usefulness of the aver-

age). When the "cautious maximin" (in which the player's interest may

want to know, "I can't make a better choice, which is not the

worst," the "optimist" may prefer to say, "I can make the best outcome

on the board." The latter statement is just as certain, and may

be more so.

This property may attract some players but not others. The

"cautious" maximin, which the principle of choice of achieving a minimum outcome

which is usually better than the maximin outcome.

Along with the certainty that the worst possible outcome is the maximin

outcome, the certainty that the best possible outcome may be the maximin

outcome is a particular advantage. The very highest element in the row may not

be such a maximum the maximin element, and it may be very low compared to

This assumption is NOT made in any game or analysis

8.

possible outcomes under other strategies. Depending on the temperament of the player and the structure of the particular matrix, the latter certainty may be so distasteful as to outweigh the attraction of the former.

Von Neumann and Morgenstern give ^{closest} ~~special~~ attention to games with the special property that \min_j [?] ; this common value is known as the "saddlepoint" of the matrix, denoted $\min_j \max_i a_{ij}$. In other words, in games with a saddlepoint, the greatest of the row-minima equals the least of the column maxima (though in the general case, the maximin element will always be less than, when not equal to, the minimax). In such games, if a player expected with certainty that his opponent was going to play his maximin strategy, the only reasonable choice for him would be his own maximin strategy, since any other would certainly give him an inferior strategy. This point is not obviously relevant, since our basic assumption is that a player never does know his opponent's strategy with certainty.

However, the authors proceed to build a "rudimentary dynamic" argument on this property. If both players, for any reason, did use their maximin strategies, the outcome of the play would be the saddlepoint value. It follows that in a subsequent sequence of plays (to take a dynamic point of view), if each player expects his opponent to continue using the same (maximin) strategy, then neither will have any incentive to change his own (maximin) strategy; so the saddlepoint value will persist. By contrast, if the matrix has no saddlepoint, then if each expects (with certainty) that his opponent will (continue to) use his maximin strategy, each will see an advantage in choosing (changing to) a non-maximin strategy. Thus, the saddlepoint, when it exists, would seem to have the properties of

both revealed choices, or each found the other out, ...
fancy words

a stationary equilibrium solution, given the assumption on expectations. In fact, in zero-sum games without a saddlepoint, there would be no equilibrium solution of this sort, under the same special assumption about expectations; whatever strategies the players chose, they would change ~~gx~~ in the next play if they ^{felt} ~~xxxx~~ certain that their opponents would not change.

But how stable is this equilibrium? Would the expectation that the opponent would use his maximin ~~x~~strategy establish itself from the beginning, or would the players begin out of equilibrium; if the latter, why and when would the expectation develop? ~~After this~~ ~~saddlepoint~~ Would the expectation persist although the opponent should in fact choose non-maximin strategies? The practical importance of this "equilibrium" depends on whether it would be attained, and whether it would be restored if one or both players should be displaced from it. These questions are crucial even to a "rudimentary" dynamic argument, and von Neumann and Morgenstern do not ~~begin~~ begin to consider them.

The limited usefulness of such an "equilibrium solution" will be apparent if it is related to the numerous "equilibrium solutions" which have been proposed for duopoly theory. This particular solution is not immediately relevant to duopoly problems, for those almost always assume non-zero-sum "games," but "solutions" to those problems always make similar assumptions about the player's expectations of his opponent's strategy. Neither the expectation that the opponent will use a minimax strategy, nor the expectation that the opponent will continue to use any particular strategy, seems ~~any more~~ uniquely reasonable ~~than~~ any more than any of the others which have been proposed to make the ~~xxx~~ duopoly problem "determinate."

The attempt to justify the significance of the saddlepoint and the minimax strategies on dynamic grounds must therefore be rejected as inconclusive, at best. At any rate, one may insist on a static approach, since many important games are played only once. A final point on the static analysis is that von Neumann and Morgenstern lend a bias to the discussion by suggesting that an "optimum" solution ~~ratherxxxprejudice~~ must guarantee that the outcome will be in some sense "optimum,"

"no matter what the opponent does." However plausible this may appear,

on reflection this whole point of view may be questioned. A reasonable person may not require a guarantee that the outcome, which will depend on the opponent's action, will be optimum in any sense. He may make his choice by balancing unfavorable outcomes against possible favorable outcomes, ~~whether or not he can form judgments about the~~

"likelihood" of different outcomes. Orthodox notions of rationality would suggest that he should then choose the strategy whose set of outcomes he most prefers. This would leave unsettled the nature of his preferences among sets of outcomes, but in any case it would ~~makexthexstxtxafxrxxtixnnality~~ define the "optimum" strategy without reference to the actual outcome which might result.

We can illustrate this point of view, and the earlier criticisms, with the following matrix.

The opt. strategy ~~to~~ should take into account all possible moves by the opponent.

$$\Delta = \begin{pmatrix} 10 & 0 & -10 \\ 0 & 0 & 0 \\ -10 & 0 & 10 \end{pmatrix}$$

The attempt to justify the identification of the...
The various strategies on dynamic grounds must therefore be rejected
as inconclusive, at best. At any rate, one has to be on a state
dependent, since many important cases are played out in a final
point on the state analysis in that von Neumann and Morgenstern
lead a line to the discussion of "optimal" solutions
must guarantee that the players will be in some sense "optimal".
no matter what the players "choose". However, this is not a
on reflection that whole point of the analysis. A player's
role period may not require a guarantee that the players, who will
depend on the opponent's action, will be rational. It may be
make his choice by choosing strategies which are rational in the
favorable outcome, exactly as in the case of the "optimal" solution.
"rational" or different classes. Opponent's actions of rationality
would suggest that he would then choose a strategy which would
outcome he most prefers. This would involve a calculation of
his preference over sets of outcomes. In the case of a player's
reference to the other players with a "best result".
We can illustrate this point of view, and the earlier criticism
with the following example.

Two just
crush
V=0
V=0
 $\vec{q}_1 = (0, 1, 0)$
 $\vec{q}_2 = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$

9.

in a game with a saddlepoint, if a player knew for certain in advance that his opponent was using a minimax strategy, the only reasonable choice for him would be his own minimax strategy, since any other would certainly give him an inferior outcome. This point has no immediately obvious relevance, since the player does not know for certain his opponent's strategy: *we assume that* ~~make the expectation that the~~ *opponent will play his*

These criticisms may be illustrated by an example. Consider the matrix: *minimax strategy seems no more "rational" uniquely than any of the others which have been proposed "to make the doubly problem determinate."*

correspondance principle.

Def. of optimum strategy as "the one which guarantees an outcome which is "optimum" in some sense. Actually person may not require such a guarantee.

According to von Neumann and Morgenstern, "the rational" way for A to play this game is to choose strategy A-3; B should choose B-3. This game is "strictly determined" since a saddlepoint exists: a_{1j}/a_{1j} Max 0. By playing "rationally" each player can avoid loss and keep his opponent from making any gain. Any other choice would expose the player to the chance of losing 10.

Would any all, or even most, people "otherwise reasonable" reject any other choice of strategy if given opportunity for deliberation? Suppose that A were to play a non-maximin strategy. He knows that he will do exactly as well if B plays his minimax strategy B-3, as if he had used his maximin strategy A-3. If ~~there~~ ~~is~~ ~~any~~ ~~possibility~~ ~~that~~ B is not certain to use B-3, one outcome seems as likely as the other then A stands to win or lose 10, and/since there is no apparent way in which B could be sure of anticipating A's choice, A might reasonably prefer this uncertainty to the certainty of winning 0.

FALSE

10.

A similar argument holds for B. In this game both might use non-minimax strategies even though each knew his opponent to be rational ~~inxxxx~~ under certainty and to be fully informed about the payoff matrix.

There is no way for B to be sure of "punishing" A for using a non-minimax strategy; in fact, to have a chance of inflicting any loss on A he must use a "non-rational" strategy himself. So long as A were convinced that he was not giving himself away, he could simply ascribe any loss to bad luck. Indeed, the very fact that B was found to be playing a non-minimax strategy might encourage A to think that he might "just as easily" have ~~xxxxxxx~~ won 10; and this is just the sort of gamble which, by hypothesis, looks good to A.

If there is reason to believe that the players will not reach the saddlepoint on the first play, there is just as good reason to believe that they will not "tend" toward it in successive plays. If one of them should prefer security, he cannot "punish" the other into a like policy. If they should both find themselves in the saddlepoint, presumeably that is where they wanted to be, for the moment; but if one or both should tire of the quiet life in subsequent plays, there is nothing to prevent their wandering away from it. In fact, in this particular game the saddlepoint does not seem to have any peculiar significance at all.

This example should be distinguished from another which appears similar. Savage has cited the matrix:

If A were not positive that B would play B-2, it would seem to show a taste for security bordering on the ~~xxxxxxxxxxxxxx~~ irrational (i.e., almost nullifying the hypothesis of maximizing behavior) for

11.

A to pick his maximin strategy A-2.

Carl Kaysen has presented a similar game-matrix in which every non-minimax strategy offers great potential gains and small potential losses as compared to the minimax strategy. There examples appear to make a better case for the use of a non-minimax strategy ~~offering equally high potential gains and losses~~ than the one offered first (which is symmetrical, a non-minimax strategy involving equally high potential gains and losses). However, in games with saddlepoints, the expectation that the opponent will use a minimax strategy makes it uniquely rational also to use a minimax strategy;¹

1. Since ~~xx~~ in this case the player is virtually acting under certainty, and his own minimax strategy offers the highest outcome.

and these unbalanced matrices, heavily favoring one player, create the presumption that if the opponent is reasonable and informed about the matrix (Kaysen drops these assumptions) he will, in fact, use his minimax strategy.

In fact, in Savage's example, the opponent B has only one "admissible" strategy, B-2, since B-2 dominates B-1. Therefore, if A were certain that B was informed and rational in the sense of ignoring inadmissible strategies (these are special assumptions, but rather weak ones), it would be irrational for A to play any other strategy but A-2.¹

By contrast, the game examined first is "indeterminate"--i.e., it seems plausible that neither player will choose a minimax strategy--even though each player is certain that his ~~xxx~~ opponent is reasonable and informed.

The point made earlier, that the minimax strategy would alone be reasonable if the opponent were known to be using a minimax strategy² (in games with saddlepoint) may be restated: the minimax principle would be uniquely rational²

to the minimum strategy 1-2.

Carl Haysen has presented a similar game-matrix in which every

non-minimax strategy offers great potential for an and small no-

partial losses as compared to the minimax strategy. There examples

appear to make a better case for the use of a non-minimax strategy

than the one

offered first (which is symmetrical, a non-minimax strategy involving

equally high potential gains and losses). However, in games with

asymmetries, the expectation that the opponent will use a minimax

strategy makes it uniquely rational also to use a minimax strategy;

1. Since in this case the player is virtually acting under cer-
tainly, and the own minimax strategy offers the highest outcome.

and these unbalanced matrices, heavily favoring one player, create

the presumption that the opponent is rational and informed

about the matrix (Haysen does not say this). He will, in fact,

use the minimax strategy.

In fact, a player's expectation that the opponent is rational and informed

about the matrix (Haysen does not say this). He will, in fact,

use the minimax strategy.

ignoring limitations strategies in fact is a rational in the sense of

but rather weak one, it would be rational for a player to

other strategy but 1-2.

By contrast, the game examined first is "infinitely"

i.e. it seems plausible that a player who chooses a minimax

strategy--even though each player is certain that his opponent

is rational and informed.

The point made earlier, that the player who chooses a minimax

be reasonable in the opponent's known to be rational.

strategy in games with asymmetries may be justified.

of which would be uniquely rational.

Handwritten notes in the left margin include:
- "actually in long run"
- "the Bell machine of a penk"
- "play flipping advantage of a penk"
- "Penny flipper"
- "does take by"
- "potential player"
- "it can be beaten"
- "guesser rational player"
- "a fool"

if the player were certain that his opponent was "rational" in the von Neumann-Morgenstern sense. But this is the type of assumption which the authors explicitly rejected in their introduction:

"If the superiority of 'rational behavior' over any other kind is to be established, then its description must include rules of conduct for all conceivable situations--including those where 'the others' behave irrationally, in the sense of the standards which the theory will set for them." (p. 32.)

We have seen that the rules of behavior they prescribe might not seem superior to others if it were not certain that the opponent would not follow them. Does not their theory fail their own criterion? They suggest an answer to this:

"It is possible to argue that in a zero-sum two-person game the rationality of the opponent can be assumed, because the irrationality of his opponent can never harm a player. Indeed, since there are only two players and since the sum is zero, every loss which the opponent--irrationally--inflicts upon himself, necessarily causes an equal gain to the other player." (p. 128)

We must insist that a rule does "harm" a player if it forces him to reject a set of uncertain outcomes (corresponding to an "irrational" strategy which he prefers to the set favored by the rules: unless it can be argued convincingly that his preferences are in some sense "irrational."¹ The very fact that the authors discuss the

1. To illustrate this, we have accepted the rule that it is irrational to prefer a set every element of which is inferior to every element of another set.

possibility that an opponent will violate any given set of rules indicates that any element in the whole matrix is possible. If there is a chance that the opponent will be "irrational," why not help him to inflict a large loss upon himself? The prospective pleasure of teaching the foolish player a lasting lesson might be worth incurring the risk of a small loss oneself. To rule out such pleasure is to put heavy restrictions on the player's "permissible" preferences among sets of uncertain outcomes.

game 2.

13.

To identify the von Neumann-Morgenstern brand of conservatism ~~xxxx~~ with "rationality" has two striking implications for "rational" preferences. First, it implies the postulate that in his preference ordering of sets of uncertain outcomes, the "rational" individual (prefers one to another) ranks the sets/strictly according to the least element in each: i.e., that the player invariably ranks strategies ~~xxxxxx~~ according to the least amount he might win under each. We would not call a man irrational for having preferences like these; but would we care to call a man irrational who did not? The assumption^A that B will succeed in enforcing in enforcing the lowest outcome in any row that A might select is dictated neither by the rules of the game, nor by B's state of information (which reflects uncertainty), nor by B's hostility; B would have to be gifted with extra-sensory perception to achieve this feat.¹ To act "as if" B were so gifted is the policy of the

1. If B merely played his minimax strategy, the result would not in general be the lowest element on the row unless A had played his maximin strategy.

"cautious pessimist": reasonable, but not uniquely so.

The second consequence is that in a game with a saddlepoint, it is the ordering of the elements in the payoff matrix that is alone relevant to choice, not their cardinal magnitudes. In other words, rational choice, as vonNeumann and Morgenstern define it, is unaffected if a payoff function (whose matrix has a saddlepoint) is replaced by a matrix which is related to the first by any increasing monotonic transformation. So long as the matrix has a saddlepoint² it is entirely

2. Two conditions are necessary involving concepts which have not been discussed; the matrix must be "specially strictly determined," i.e., the saddlepoint must correspond to a pair of "pure" strategies, and second, there must be no chance moves in the extensive form of the game.

unnecessary that the payoff be expressed in money or any cardinal magnitude; any index expressing the player's ordinal preferences

To identify the von Neumann-Morgenstern basis of rationality with "rationality" has two striking implications for "rationality". First, it implies the postulate that in his preference ordering of acts of uncertain outcome, the "rational" individual ranks the acts/alternatives according to the least element in each: i.e., that the player invariably ranks alternatives according to the least amount he might win under each. We would not call a man irrational for having preferences like these; but would we call a man irrational who did not? The assumption that B will succeed in enforcing in enforcing the lowest outcome in any row that might select is dictated neither by the rules of the game, by A's state of information (which reflects uncertainty), nor by B's rationality; B would have to be gifted with extraordinary perception to achieve this feat. To see "as if" B were so gifted is the policy of the I. If B merely played his maximin strategy, he would not in general be the lowest element in any row unless he had the maximin strategy.

This is of course psychologically biased.

on the difference between

Utilities and \$

This comment however is in passing - the Russians

en passant - is more

counterexample is important

necessary that the payoff be expressed in money or any cardinal magnitude; any index expressing the player's ordinal preferences

14.

among outcomes (not set of outcomes) would suffice, so far as the prediction of their choice among strategies is concerned.²

This property deserves ~~xxx~~ a good deal of thought. It seems particularly unrealistic to assume that the behavior of all or most reasonable people would be unaffected by a monotonic transformation of the payoff function (e.g., replacing each outcome by its square, ^{if each all were positive} in the same units). Surely many people would be interested in comparing the differential gains that they might make by choosing a non-minimax strategy to the differential losses they would risk. In the first matrix cited, a player who was willing to accept the uncertainty of receiving either 10¢, 0¢, or -10¢ might be unwilling to risk the loss of \$100, even if combined with possibility of winning \$100.¹ Yet this transformation would be

1. It is no solution to imagine that the outcomes are expressed in "von Neumann-Morgenstern utilities" (and at any rate, they are not), since those are employed only to formalize choice in situations involving "risk," i.e., where a probability distribution is known.

entirely disregarded by a player who was rational according to the von Neumann-Morgenstern rules, since he would have chosen the minimax strategy in the first place.

money v. utility

In The Theory of Games, the solution which von Neumann and Morgenstern propose appears much more plausible than it does in the above discussion, because it is presented first in connection with some modifications of the game-model, then applied/without close examination to the normalized game. It will be argued below that there are crucial ~~differe~~ features of ~~between~~ the modified forms, which are supposedly introduced merely for ~~didactic~~ pedagogic reasons, which make extrapolation to the normalized form invalid.

In the normalized ~~form~~ game, which is the primary subject for analysis, both players choose strategies simultaneously, each in ignorance of the other's choice. In the first modified game, called the minorant game, A must make his choice first, after which B chooses in full knowledge of A's choice. Since B, in this game, acts under certainty, the basic principles of rationality under certainty prescribe his choice. Given strategy i by A, B's unique rational choice is that strategy which minimizes a_{ij} ; i.e., he should pick the column corresponding to the minimum element in the row selected by A. Given A's strategy, then to each strategy available to B there corresponds a single, certain outcome, and rationality compels him to pick the strategy associated with the outcome: $\text{Min}_j a_{ij}$, where i is given.¹

^{in that case}
1. Strictly speaking, this strategy might not be unique; but/the ~~unique~~ minimum value $\text{Min}_j a_{ij}$ would be the same for all the "rationally admissible" strategies.^{1j}

A's problem is not quite so simple, but it can be made so by a relatively weak assumption. If A does not know B to be rational--a fortiori, if he knows that B is not rational--then A must choose under some degree of uncertainty. But if A knows, for certain, that B is rational under certainty (this knowledge is our special assumption), then A, too, acts under certainty.

that the choice of idea forced with a set the choice accept

individual of choices back you your the game the min

maximal one by the game the min

this induction up the game the min

accept game the min

In The Theory of Games, the solution which was chosen and
Korotkiy proposes suggests that the solution which was chosen in the
above discussion, because this procedure is based on the assumption
some additional of the game-models, then applied to the normalized
game. It will be shown below that these are the same as the
between the modified forms, which are usually introduced
for strategic reasons, which make extrapolation to the
normalized form invalid.

In the normalized form game, which is the ordinary subject for
analysis, both players choose strategies simultaneously, and the
ignorance of the other's choice. In the first model called
the normal game, a player chooses first, after which
chooses in full knowledge of the other's choice. In this game,
sole under certainty, the basic principle is that a player
certainly prescribe his strategy. Give strategy 1 by A's unique
rational choice is that strategy which minimizes his loss; he should
pick the column corresponding to the minimal element in the row
selected by A. Given A's strategy, then to each strategy available
to B there corresponds a single, certain outcome, and accordingly
compels him to pick the strategy associated with the outcome Min.
A, where 1 is given.

I. Strictly speaking, this strategy might not be unique; but the
outcome Min. would be the same for all the rational
ally rational strategies.

A's problem is not quite so simple, but it is made so by
a relatively weak assumption. If A does not know B to be rational--
a fortiori, if he knows that B is not rational--then A must choose
under some degree of uncertainty. But if A knows, for certain,
that B is rational under certainty, then knowledge is not essential
assumption, then A, too, can get certainty.

If A knows (believes with certainty) that B is rational under certainty, then in the minorant game (as contrasted with the normalized game) this belief has definite implications for A's expectations and behavior. Operationally, the statements that B is ^{acting under certainty,} certainly rational and that B, is certain to choose the minimum element in any row picked by A are identical. The knowledge that B is rational then implies that it is impossible that a given strategy by A should have as outcome any element but the minimum in the row; hence, it is irrational for A to pay any attention to the $m(n-1)$ matrix elements that are not row minima.¹ This leaves A with m

1. The postulate that it is irrational for a player to be influenced in his choice of strategy by outcomes which he considers ~~absolutely~~ impossible is somewhat different from those accepted earlier, but surely it is equally acceptable.

possible outcomes, one corresponding to each of his m strategies.

STEP

~~Obviously, he should choose the "maximin" strategy corresponding to the largest ~~of~~ row minimum.~~

① As stated earlier, a solution which depends on the assumption that one player has special knowledge about the other is not acceptable as a general solution, even to this special game. On the other hand, it should be noticed that the assumption made here is merely that A believes B to be rational under certainty, a concept that is well defined; no use is made of any concept of rationality under uncertainty. It is certainly significant that in this minorant game this limited assumption makes the outcome determinate, whereas it is irrelevant to the normalized game.

The second special model is called the Majorant game: in this, B must choose before A, who then makes his choice in full knowledge of B's choice. Now A chooses with certainty. ^{As above,} /If B knows A to be rational ~~under~~ certainty, this is equivalent to knowing that elements which are not column maxima are not possible outcomes. Hence under

17.

this special assumption B also acts under certainty, associating a single, certain outcome ($\text{Max}_i a_{ij}$, for given j) with each strategy. B's only strategy which is rationally consistent with his assumption about A is his minimax strategy, which guarantees him the ~~ix~~ best "possible" outcome.

~~Although we conclude that in these special games the first~~

Although we conclude that in these special games the first player should choose the strategy which happens to be "rational" in the von Neumann-Morgenstern sense, provided that he assumes the second player to be rational under certainty, the reasoning behind this conclusion is obviously ~~xxxxxx~~ peculiar to these special games. It may be put into terms familiar from ~~xxxxxx~~ duopoly theory.

Since the second player need not take reactions into account, in these particular/^{static}models, rational choice is for him a clear-cut matter; therefore it is possible to prescribe for him a "reaction function" which is consistent with the assumption of his rationality. If the first player assumes that the ^{second} ~~first~~ player is rational in this broad sense he can deduce his opponent's reaction function, and his task is merely to ~~xxx~~ pick the point ~~xxx~~ on that function most favorable to himself.

The discussion by vonNeumann and Morgenstern is faulty in an important respect; they do not assume explicitly that the first player knows the second to be rational in any sense.¹ Yet, in ~~xxxxxxxxxxxx~~

1. Nor that the first knows the second to be informed about the ~~FALSE~~ payoff matrix. They state explicitly in the introductory discussion (p. 30) ~~that they assume all players are fully informed; but much of their argument suffers from lack of an explicit postulate that all players make similar xxx assumptions about each other.~~

discussing the minorant game, they state unqualifiedly that B is "certain" to minimize a_{ij} for any given i , and that A knows this:

hence that when A picks a particular strategy "he can already fore-see with certainty" what his outcome will be. (p. 101). They ^{fail to} arrive ~~at the conclusion that A acts under certainty without assuming explicitly A's knowledge of B's rationality.~~ *explicit*

~~Nevertheless, this assumption is implicit.~~
~~This procedure is invalid.~~ Without ~~xxx~~ certainty that B is rational and informed, there can be no "certainty" of outcome for A. If A were not sure that B was rational, it would not be irrational ~~strict~~ traditional (in the ~~broad~~ sense) for him to pay attention to other outcomes than row minima. So long as there was a genuine possibility that he might attain them (i.e., that B might choose an outcome less than he could achieve with certainty, either from non-rational motives or from ignorance of the payoff) he might "reasonably" be attracted to a non-maximin strategy by hopes of large gains. In a more dynamic analysis, B could pursue a strategy of a sort that von Neumann and Morgenstern never consider: luring A away from a maximin policy by creating doubts as to his own rationality (e.g., by taking ^{greater} ~~xxx~~ than the minimum element in the maximin row).

Even though uncertainty should exist in the mind of the first player, it might be argued that he should pursue the maximin (or minimax) policy anyway, since this would have a better consequence than any other if the opponent should prove rational. It was argued above that this principle would not have a unique claim to reasonableness. At any rate, in this context it is not the one that von Neumann and Morgenstern propose. They really make the key assumption implicitly, that A is certain (that B is rational), rather than argue that A should act "as though" he were certain; this is clear from the quotations above. It is also implied by their motives in discussing the games:

See p. 100

"The introduction of these two games...achieves this: it ought to be evident by common sense--and we shall also establish it by an exact ~~argument~~ discussion--that for (these games) the 'best way of playing'--i.e. the concept of rational behavior--has a clear meaning." (p. 100).

Our discussion of these special games supports their conclusion that (granted the assumptions implicit in their analysis), "The good way" (my italics) for each player to play these respective games can be prescribed. (pp. 101-103, in particular paragraphs 14:A:a-14:A:e and 14:B:a--14:B:e). ¹⁰³

The essential fact about these special games is that the players' beliefs about his opponent's rationality under certainty can remove all uncertainty from his own choice-situation. It is this very fact, which makes the special games interesting in themselves, which makes them basically different from the normalized game, in which it is impossible to banish uncertainty by any such simple assumption, and in which uncertainty is the essence of the problem. This makes any attempt to ~~extrapolate~~ apply the results of their analysis to that of the normalized game suspect from the beginning. It so happens that it is possible to locate the exact spot where von Neumann and Morgenstern hurtle the gap.

Since make outcome "determinate." In terms of monopoly theory. ①
although this is not acceptable... ②

all your criticisms can be directed against the O₁ of individual consumer demand - without which there is more or less no micro - ec O₄ whatever - Hicks, Samuelson etc must be then thrown away

In their discussion of the normalized game, von Neumann and Morgenstern begin by ~~maximin~~ discussing the general advantages (from a conservative point of view) of the minimax principle. Then, in section 14.5, they approach for the first time directly the definition of rational choice in the normalized game. They start out: "It is reasonable to define a good way for 1 to play the game" as that strategy that will ~~give~~ guarantee him at least the maximin outcome, ~~X~~ Similarly, "it is reasonable to define a good way for 2 to play the game as one which guarantees him a gain" ~~which~~ corresponding to the minimax outcome.

So far there can be no quarrel with these statements; the principle they describe (which, incidentally, is an old one) cannot surely be called unreasonable. They continue:

"So we have:

(14:C:a) "The good way (strategy) for 1 to play the game" is ~~the~~ maximin.

And:

(14:C:b) "The good way (strategy) for 2 to play the game" is minimax. They conclude, at the bottom of the page:

"Finally, our definition of the good way of playing, as stated at the beginning of this section, yields immediately...." (p. 108)

The **Fact** is that ~~their~~ statement at the beginning of the section did not define "the" good way of playing. It defined "a" good way. It was not until three paragraphs later that "~~a~~" was quietly transmuted into "the" (the italics above, of course, are mine). Nevertheless, the authors feel free to start ~~the~~ next section with the statement: "(14:C:a)-(14:C:b)...settle everything as far as the strictly determined two-person games are concerned." (p. 109; "strictly determined" means that the game has a saddlepoint). For this class of strictly deter-

is then discussed of the normalized game, von Neumann and
Morgenstern begin by discussing the general advantages
(from a conservative point of view) of the minimax principle. Then,
in section 1A.5, they introduce for the first time directly the
definition of rational choice in the normalized game. They state that
"it is reasonable to define a good way for I to play the game" as that
strategy that will give guarantee at least the maximum outcome.
Similarly, "it is reasonable to define a good way for S to play the
game as one which guarantees at least a gain" which corresponds to the
minimax strategy.
So far there can be no conflict with these statements of the principle
they describe (which, incidentally, if one cannot strictly be
called unproblematic. They continue:
"So we have:
(1A:3:a) "The good way (strategy) for I to play the game" is the
maximin.
and:
(1A:3:b) "The good way (strategy) for S to play the game" is minimax.
They conclude, at the bottom of the page:
"Finally, our definition of the good way of playing, as stated
at the beginning of this section, is immediately..." (p. 108)
The fact is that their statement at the beginning of the section did
not define "the good way of playing. It defined "a" good way. It was
not until three paragraphs later that "a" was directly translated
into "the" (the ideas above, of course, are clear. Nevertheless,
the authors feel free to state the next section with the statement:
"(1A:3:a)-(1A:3:b)...we have everything as far as the strictly determined
two-person games are concerned." (p. 109; "strictly determined" means
that the game has a saddlepoint). For this class of strictly deter-

unique solution
a
may have many
strategies in it

games they are now satisfied that they have "a precise theory..which gives complete answers to all questions." (. 101).

The metamorphosis of "a" into "the" on this page is no mere printer's error, ~~and~~ nothing that can be "clarified" by a footnote in later editions. The whole structure of their "determinate" theory (including that of the n-person game, which requires a unique solution to the two-person game), the whole of their claim to have recognized the true stature of a timeworn maxim, rests on a basis no more substantial than this.

It is the keystone of the whole structure of their "determinate" theory (including that of the n-person game, which requires a unique solution to the two-person game). The alchemist's magic which transmutes a timeworn maxim into an overriding postulate of rational choice is, after all, a bit of sleight-of-hand.

The limitations of the von Neumann-Morgenstern analysis can be firmly established in terms of their approach to the type of games discussed above. However, a thorough discussion must consider what the authors ~~present~~ themselves emphasize much more: the significance of von Neumann's theorem concerning the existence of saddlepoints.

Once again, certain key concepts are introduced in connection with the minorant and majorant games. In the ~~first~~ former, it will be recalled, if A (moving first) ^{believes} ~~knows~~ B to be rational under certainty A will choose his maximin strategy; if B is in fact rational, he will choose the column corresponding to the minimum element in the row chose by A, so the outcome is uniquely determined: $\text{Max}_i \text{Min}_j a_{ij}$. In the majorant game, if A (moving second) is rational under certainty, and if B knows this, then B (moving first) will pick his minimax strategy and the outcome of the game will be: $\text{Max}_j \text{Min}_i \text{Max}_i a_{ij}$.

do you claim that this assumption is bad?

It is the peculiarity of these special models that the special assumptions that both players are rational under certainty and that the player moving first knows this of his opponent make the outcome of the play uniquely determined. ~~Under these special conditions~~ So long as these special conditions apply it is plausible to define the unique outcomes as the "values" of the respective games for the two players. In the minorant game, the outcome to A under the special conditions will be $v_1 = \max_i \min_j a_{ij}$ and the outcome to B will be $-v_1$. In the majorant game the outcome to A is $v_2 = \min_j \max_i a_{ij}$, the outcome to B $-v_2$.

It can be proven mathematically that v_1 (maximin) is always less than or equal to v_2 (minimax). A would always prefer to play the majorant than the minorant game: if the matrix ~~was~~ ^{was} the same for each, B was known to be rational under certainty, and the game was zero-sum. ((It is true that the minorant game is "clearly" "less advantageous for A than the majorant game (p. 100 and p. 105)

((If the game were not zero-sum, it would no longer be true that the minorant game was "clearly" "less advantageous" for A than the majorant game. ~~Ques~~ In a non-zero-sum game, in which it is possible for both players to "lose" simultaneously, one player may wish strongly to let his opponent know his intentions, preferably by moving first. This situation is discussed in an unpublished doctoral thesis by Howard Raiffa of the University of Michigan.

Given these conditions, it would be possible to prescribe a "rational" preference ordering of games (i.e. of payoff matrices) for either player, in terms of the relative "values"/for each. (as defined above)

If $v_1 = v_2$, a "saddlepoint" is said to exist, denoted by: $S_{i_1/j_1} a_{i_1 j_1}$. The existence or non-existence of a saddlepoint is of no interest at all in the minorant and majorant games so far as the behavior of the players is concerned. The only behavioristic significance it would have would be that a player would be indifferent between playing the minorant or the majorant games if he knew his opponent to be rational under certainty.

It is the possibility of these special moves that the special

assumption that both players are rational under certainty and that

the player moving first knows this of his opponent make the outcome
to look as these special conditions
of the play uniquely determined.

It is possible to define the unique outcome as the "value" of the

respective games for the two players. In the minorant game, the out-

come to A under the special condition will be v_1 and in

outcome to B will be $-v_1$. In the majorant game the outcome to A is

v_2 and to B $-v_2$.

It can be proved mathematically that v_1 (maximin) is always

less than or equal to v_2 (minimax). A would always prefer to play

the majorant than the minorant game: if the matrix were the same for

each, B was known to be rational under certainty, and the game was

zero-sum. // It is true that the minorant game is "clearly" "less

advantageous for A than the majorant game (0.500 and 0.100)

((If the game were not zero-sum, it is no longer true that the

minorant game was "clearly" "less advantageous" for A than the

majorant game. But in a non-zero-sum game, in which it is possible

for both players to "lose" simultaneously, one player may wish

attempts to let his opponent know his intentions, thereby

moving first. This situation is discussed in an unpublished report

of Howard Raiffa of the University of Michigan.

Given these conditions, it would be desirable to prescribe a "rational"

"preference ordering of games" (a. of p. 100) (Raiffa)

player, in terms of the relative "value" for each.

If $v_1 = v_2$, a "saddlepoint" is said to exist, and we

say v . The existence or non-existence of a saddlepoint is of no

interest at all in the minorant and majorant games so far as the

behavior of the players is concerned. The only behavioral sig-

nificance it could have would be that a player would be indifferent

between playing the minorant or the majorant game if he knew his

opponent to be rational under certainty.

Handwritten notes:
significantly
upper
they are
lower
and points

At this point in their discussion, the authors point out: "Thus far we have not even attempted the proof that a numerical value of a play can be defined in this manner" for the normalized game. (p. 105). Such a proof would seem to require that some relatively weak assumptions, such as that both players are informed about the matrix, and perhaps that both players know ~~each~~ each other to be rational under certainty ((admittedly, this assumption is ^{with} "weak" only by comparison ~~by~~ some of those which von Neumann and Morgenstern see fit to make)). logically imply a unique outcomes for both players.

~~They begin~~ The essence of their proof appears in their initial statement: "Instead of ascribing v_1 , v_2 as values to ~~the normalized game~~ ~~and~~ ...two games different from ((the normalized game))~~x~~ we may alternatively correlate them with ((the normalized game)) itself." (p. 105). This ~~argument~~ line of approach seems to be on treacherous footing from the start. They proceed with a "heuristic" argument to suggest that the numbers v_1 and v_2 have a practical significance in connection with the normalized game. Although in this game both players choose simultaneously, "It is nevertheless conceivable that one of the players, say 2, 'finds out' his adversary; i.e., that he has somehow acquired the knowledge as to what his adversary's strategy is." They assert that in this case, conditions "become exactly the same as if the game were" the minorant game. Likewise, if player 1 "finds out" his adversary, conditions become "exactly the same as it" the game were the majorant game. Hence they claim that in either of these cases the "value" of the normalized game becomes a "well-defined quantity": v_1 in the first case, v_2 in the second. (p. 106) In general, the implication might be drawn that v_1

24.

is the ^{maximum} amount that player 2 "should" be willing~~x~~ to pay for privilege of playing the game, ^{if he were} with the guarantee~~d~~ of foreknowledge of 1's strategy.

There is a basic flaw in this argument. Before discussing the ~~xxxxxxxxxxxx~~ consequences of the possibility that A will "find out" B, we must assume that B would not be aware, in advance of picking his strategy, that he was to be found out. This assumption is in conflict with the authors' conclusions, but in harmony with their "heuristic" argument. After all, if B knew for certain that he would be found out, then it would not be "as if" they were playing the majorant game; they would be playing the majorant game.

If this assumption is granted, then the inference to be drawn from the whole of our previous discussion is that B might "reasonably" ^(found) be playing some non-minimax strategy. The reward to A of "finding out" B under these circumstances is not limited to v_2 (minimax), though that is a lower bound; ~~xxxxxxxmightxxxxxxxachievexxxxxxxthehighest~~ with foreknowledge in the normalized game, A might be able to achieve the very highest outcome in the matrix, even though B were rational under certainty. ~~xxxxxxxmightxxxxxxxwillxxxxxxx~~ In other words, if it is accepted that an opponent may "reasonably" choose a non-minimax strategy in the normalized game, then "value of finding B out" is not limited to "the value of the majorant game" for A; it might be much more. *(which is a good reason for playing minimax)*

Similarly, the possibility that B may find out A implies that the final outcome may range anywhere from maximin down to "minimin," the lowest element in the matrix. Both possibilities together imply that the outcome may range from minimin to maximax, i.e., may take on any value in the payoff function., even though each player ~~xxxxxx~~ should know the other to be rational under certainty ((~~this unneces-~~ ^{irrelevant line,} ~~sary~~ assumption is mentioned only to contrast the situation with that

The way you have treated it, makes this true

maximum
in the amount that player 2 "should" be willing to pay for privilege
of playing the game with the guarantee of knowledge of 1's
strategy.

There is a basic flaw in this argument. Before discussing the
game-theoretic consequences of the possibility that 2 will "find out"
it, we must assume that 2 would not be aware, in advance of playing
his strategy, that he was to be 2 and out. This assumption is in
conflict with the authors' conclusions, but in harmony with their
"naive" argument. After all, if 2 knew for certain that he
would be found out, then it would not be "as if" they were playing
the relevant game; they would be playing the relevant game.
If this assumption is granted, then the inference to be drawn
from the whole of our previous discussion is that 2 might "be naive"
by playing some non-minimax strategy. The reward to 2 of "finding
out" 2 under these circumstances is not limited to v_2 (minimax),
though that is a lower bound; ~~xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx~~ with
foreknowledge in the relevant game, 2 might be able to achieve the
very highest outcome in the matrix, even though 2 were not that under
certainly. ~~xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx~~ In other words, if it is
accepted that an opponent may "be naive" ~~xxxxxxxxxxxxxxxxxxxxxxxxxxxxxxxx~~
strategy in the non-relevant game, then "value of finding 2 out" is
not limited to "the value of the relevant game" for 2; it might be
much more.

Similarly, the possibility that 2 find out 1's strategy does
the final outcome may range anywhere from maximin down to "minimin."
The lowest element in the matrix, not possibilities together imply
that the outcome may range from minimin to maximin, i.e., may take
on any value in the payoff function, even though each player knows
should know the other to be rational under certainty (this unques-
tioned assumption is mentioned only to contrast the situation with that

her
not

25.

minorant and majorant
of the ~~normalized~~ games, in which it has a decisive effect.))

Unless both players are ultra-conservative (~~defined~~ i.e., rely exclusively on the minimax principle) there seems no warrant for restricting the range of outcomes that might result if one player should find out the other to v_1 -- v_2 .

This conclusion would be ~~fat~~ damaging to the von Neumann-Morgenstern argument. After establishing to their satisfaction that the interval v_1 -- v_2 represents "the advantage" to be gained from "finding out" one's adversary instead of being "found out" by him" (p. 106) they draw the conclusion that games with a saddlepoint ($v_1 = v_2$) acquire a peculiar significance, in that "it does not matter which player 'finds out' his opponent." (p. 106). If the criticism above is valid, then this significance fades away. If v_1 and v_2 , separately have little relevance to the normalized game, they are no more relevant when they happen to ^{be} equal. ~~xxxxxx~~.

into Koyen

signif. of v_1 & v_2 .

Nevertheless, we can rescue some scraps of significance ~~xxxxx~~ for the saddlepoint concept by admitting that ~~xxxxx~~ a few restrictive statements (applying to "defensive" players) can be made about games with a saddlepoint, though it is necessary to resort to a crude dynamic argument. If a saddlepoint exists, it represents an outcome if both $v = v_1 = v_2$ such that A can be sure of receiving at least v no matter what B does, and B can keep A from receiving more than v no matter what ~~Ex~~ A does. To say that consideration of v is the sole determinant of their behavior, no matter what the ~~xx~~ remaining structure of the matrix, is to say that they are both predominantly concerned with security, with a sure minimum outcome. If this is the case, the v is the outcome which will actually result. The same result will follow if a saddlepoint exists and one player (who ~~xxxxxx~~ need not be conservative) knows the other player to be conservative.

|||||
?
2

From the static point of view this result is purely formal, for it is equally true that in any game whatever, even without a saddlepoint, the assumptions that ^(a) both players are conservative or that one ^(b) knows the other to be conservative make the outcome uniquely determinate. The only difference is that without a saddlepoint it would not be possible to characterize the outcome of all games ^{for conservative} ~~such~~ players by the single abstract symbol: $Sa_{1/j} a_{1j}$; the most that can be said abstractly and generally in such cases is that the outcome will lie between v_1 and v_2 (for ^{nonetheless} ~~such~~ players). A). For a given game, the outcome would be ~~unique and~~ predictable. Statically, then, the existence or not of a saddlepoint would seem to have no effect at all on the process of choice of inherently conservative players.

Similarly, if one player knew the other to be conservative the outcome would be ~~determined~~ determined: either v_1 or v_2 depending on which player had this information.

But in a dynamic analysis, it might be argued that in the absence of a saddlepoint even the assumptions a) or b) above would not make the outcome determinate. Under (a), the fact that ^a ~~the~~ player would realize that he could better his outcome by abandoning his minimax strategy if only his opponent retained his conservative policy might tempt him to betray his innately conservative temperament. Under (b), the traditionally conservative player might be tempted to punish his opponent for exploiting him ^{in the past.} ~~formerly~~. Neither temptation, it should be noted, would exist if the game had a saddlepoint. In the course of several plays, then, there would be some pressure for one or both players to abandon their conservatism, so that the full indeterminacy would reappear. The authors' assertion that v_1--v_2 is the significant interval of indeterminacy in the absence of a saddlepoint seems dubious; the whole range of outcomes would seem to be possible.

The discussion so far has suggested that the existence of a saddlepoint is of strictly limited significance, but not entirely without interest. Von Neumann and Morgenstern suggest an approach whereby every game matrix could be considered to have a saddlepoint.

The essential point is that a game permitting ~~xxx~~ a set of specified strategies can be considered to permit any probability combination of these strategies. If it is permissible to play strategies 1 or 2, then the rules cannot prevent a player from deciding between them by flipping a coin or rolling dice. ~~In this~~ ~~first case~~ He might then be said to be "playing both" with fixed probabilities, in the first case each strategy having probability $\frac{1}{2}$. The real choice he makes then is the rule for correlating particular strategies with probabilities; e.g., he might decide to play 1 if the coin lands heads, 2 otherwise. Or he may mark one card "1," nine cards "2", shuffle them and pick one, playing the corresponding strategy. This would be equivalent to ^{choosing} ~~playing~~ strategy 1 with probability 1/10 and strategy 2 with probability 9/10. In any case, of course, he ends up playing one particular strategy; the probabilities merely refer to the random nature of his choice. If he chooses to play one strategy with probability 1 and all others with probability 0, he is said to choose a "pure" strategy, of the sort considered exclusively until now. In general, the player can be said to choose a "mixed strategy": i.e., to choose all strategies with fixed probabilities, the basic decision being the choice of a vector of probabilities.

Associated with ~~Corresponding to~~ each "mixed" strategy there will be a set of possible probability distributions of outcomes (instead of a set of definite outcomes), each distribution corresponding to a particular pure or mixed strategy by the opponent. Von Neumann and Morgenstern

The discussion so far has suggested that the existence of a saddle point is of strictly limited significance, but of entirely without interest. Von Neumann and Morgenstern suggest an approach whereby every game could be converted into a saddle point game. The essential point is that a game is converted into a saddle point game by the addition of a constant to the payoffs. It is in this sense that the addition of a constant to the payoffs is of no importance in the decision between two strategies. This would be equivalent to mixing strategy 1 with strategy 2 with probability 0.50. In any case, of course, he ends up playing one particular strategy; the probabilities merely refer to the random aspect of his choice. It is chosen to play one strategy with probability 1 and all others with probability 0, as is said to choose a "pure" strategy, of the sort considered exclusively until now. In general, to choose all strategies with fixed probabilities, and make a decision on the basis of the expected value of probabilities.

In the whole book there is the implicit assumption for purposes of division of difficulties that money is utility

They are saying in most cases as an approach to utility (say flipping for coffee so) money is roughly

now make use of a powerful assumption which has not been discussed earlier. They assume that a player considers only the "mathematical expectation" of a distribution of outcomes: i.e., the mean of the set of outcomes weighted by their respective probabilities. In other words, they assume that a player will prefer one distribution to another if and only if the first has a higher mathematical expectation ~~than~~ than the second. Hence the set of distributions associated with a given mixed strategy can be represented by a set of numbers, the numbers being mathematical expectations. The matrix of the game will now be much larger, with new rows and columns corresponding to all the possible mixed strategies, but it will have the conventional appearance.

The basic theorem of von Neumann, first proven in 1928, is that every matrix in which mixed strategies are included will have a saddlepoint. In their terminology, ^(the logic of which has been questioned) every game is "strictly determined." The game with a saddlepoint corresponding to a pair of "pure" strategies is "specially strictly determined." It remains true that this saddlepoint cannot be significant except under the conditions noted above for the "specially strictly determined" game; moreover, it requires additional ~~assumptions~~ conditions to be empirically relevant. ~~First, the assumption that~~ First, the assumption that the players ~~order~~ rank distributions of outcome in an order of preference corresponding strictly to their mathematical expectations (paying no attention, for example, to range or variance) will be unpalatable to many economists. It should be recalled that the outcome is expressed in terms of money. ~~Von Neumann and Morgenstern~~ The discussion by von Neumann and Morgenstern in their introduction of a possible index of "utility" in terms of which most people could be said to maximize expected "utility" is irrelevant here. It is unequivocally assumed that what the players are maximizing is the mathematical expectation of money. This is recognized by the authors to be only an approximation, but it is one that ~~they cannot dispense with in their theory~~ plays a crucial

now take use of a somewhat assumption which has not been discussed
earlier. They assume that a player considers only the "mathematical
expectation" of a distribution of outcomes; i.e., the mean of the
set of outcomes weighted by their respective probabilities. In
other words, they assume that a player will not be one distribution
to another if and only if the first has a higher mathematical expect-
ation than the second. Hence the set of distributions
associated with a given fixed strategy can be represented by a set
of numbers, the numbers being mathematical expectations. The matrix
of the game will now be such larger with rows and columns corres-
ponding to all the possible mixed strategies. It will have the
same overall appearance.

The basic theorem of von Neumann and Morgenstern is that
every matrix in which mixed strategies are included will have a
saddlepoint. In their terminology, every game is "strictly determined".
The game with a saddlepoint corresponds to a set of "pure" strategies.
It remains true that this saddlepoint cannot be strictly deter-
mined as "essentially strictly determined".

Only under the conditions noted above for the "essentially strictly
determined" game; moreover, it requires additional mathematical ex-
pectations. First, the assumption that
the players rank distributions of outcomes in an order of prefer-
ence corresponding strictly to their mathematical expectations (paying
no attention, for example, to range or variance) will be unpalatable
to many economists. It should be recalled that the outcome is expressed
in terms of money. Von Neumann and Morgenstern in their introduction of a "utility
index" of "utility" in terms of which most people would be said to rank the expected
"utility" is irrelevant. It is unconvincingly assumed that what
the player the maximization is the mathematical expectation of money.
This is recognized by the authors to be only an approximation, but
it is not that the mathematical expectation of money is a principle

What is a
conservative
player?

role in the theory as developed by von Neumann and Morgenstern.

Second, the use of mixed strategies to ensure the existence of a "solution" has little intuitive appeal. Many of the very players who were conservative enough to use the minimax principle with respect to pure strategies would be uninterested in the "optimal" mixed strategy, precisely because they would still tend to consider¹ mainly the lowest possible outcome under each strategy. ((This amounts to saying that it is precisely those players who are conservative enough to use the minimax principle who would be likely to consider other aspects of a probability distribution than its expectation; they would be concerned over its minimum.)) By using any non-pure strategy, a conservative player A would always incur ~~the~~ a positive probability of suffering an outcome worse than the maximum outcome, possibly including the very worst outcome. In this sense, the use of any mixed strategy would involve some loss of security.

Several arguments have been suggested to make the concept of mixed strategies more acceptable. ~~Although~~ The first is ~~xxxxxx~~ a notion that the authors reject but that has appeared in other writing, the interpretation that the optimal mixed strategies are used to achieve ~~xxxxxxxxxxxxxxxx~~ the highest long-run distribution of outcomes over many plays consistent with security. Thus, Marschak says that by introducing the concept of the mixed strategy: "...not the value of a single play for player A but the long-run value of the game for player A is considered." This may be the most plausible explanation of the use of numerical probabilities. However, one might question whether really important games would be repeated frequently enough to make mean long-run value interesting. Von Neumann and Morgenstern themselves avoid ~~xxxxxxxxxx~~ such an openly dynamic argument. Their own "rationalization" of the use of mixed strategies is that the player thereby avoids being found out, since he himself does not know which strategy will be used on a particular play. But as they themselves

the game has a saddlepoint corresponding to a pair of pure strategies. Second, the consideration of mixed strategies is not sufficient to ensure the use of the minimax principle. If mixed strategies are considered at all, then any other rules of choice could take them into account, and (except in the case of maximax) the rules might dictate the use of a mixed strategy just as often as the minimax rule does. Thus, the question of the choice of a rule of behavior is logically independent of the ~~xxxxxx~~ question whether or not to consider mixed strategies.

Mixed strategies do ensure the existence of a bilinear form with a saddlepoint, a point which may be of interest to the theorist (though not to the individual player) in predicting the outcome of a game played by two minimaxing players. But they do not ensure that this bilinear form will be of interest even to the theorist; that depends on whether the players are not only minimaxers but ~~xxxx~~ are interested only in also ~~xxxxxxx~~ the mathematical expectation of risk-prospects, "lottery-tickets", when probabilities are known. (if the problem of transferability can be solved, this implies that a von Neumann-Morgenstern utility index can be found for each player). Since the latter assumption is often thought ~~xxx~~ more distasteful than the first, it is fortunate that the two are logically* independent.


more customary notation is

$$\begin{pmatrix} 10 & -10 \\ -10 & 10 \end{pmatrix}$$

the vectors represents

the pure prospects

not only would they, but they
do with a 50:50 probability



31.

Let us ~~imaginaxx~~ consider a game of Matching Pennies consisting of 10 moves (matches), in which each player was permitted only three strategies: all heads, all tails, or a mixed strategy with probabilities ~~xxxxxx~~ $(\frac{1}{2}, \frac{1}{2})$; in the latter case, the player would simply flip the coin at each move. This would have the matrix:

10	0	-10
0	0	0
-10	0	10

This matrix is essentially the same as the one considered first. ~~Thxxxxxxx~~ Where ~~xx~~ the mixed strategies $A = (H, T, \frac{1}{2}, \frac{1}{2})$ or $B = (H, T, \frac{1}{2}, \frac{1}{2})$ are involved, the outcome represents the mathematical expectation of a distribution of outcomes; but assuming that the player is interested only in mathematical expectation, this fact can be ignored. The earlier discussion obviously applies here fully, despite the changed interpretation of the strategies. To recapitulate briefly, it seems that many people faced with this matrix would choose a "bad" strategy, under which they might win or lose 10¢, rather than their "good" strategy, which would give them exactly 0¢. ~~xx~~ Von Neumann and Morgenstern concede some ambiguity in ~~xxxx~~ these terms "good" and "bad." As they point out in connection with an equivalent matrix, suppose that a player should play a "bad" strategy: "If the opponent played the good strategy, then the player's mistake would not matter." (p. 164). Then in what sense would the "bad" strategy be a mistake? If A should pick A-H and B should pick B-H, would A's 10¢ win measure the "badness" of his "mistake"? True, A would have exposed himself to the loss of 10¢; the question ~~xx~~ is whether he should be called irrational for doing so.

So much has been pointed out already. A new question ~~xx~~ that

Let us imagine consider a game of matching pennies consisting

of 10 moves (games), in which each player was permitted only three strategies: all heads, all tails, or a mixed strategy with probabilities $\frac{1}{2}$, $\frac{1}{2}$; in the latter case, the player would simply flip the coin at each move. This would have the matrix:

This matrix is essentially the same as the one considered first.

Whereas the previous matrix was $A = (H, T, \frac{1}{2}, \frac{1}{2})$ or $B = (H, T, \frac{1}{2}, \frac{1}{2})$ and the outcome represents the expected value of the player's action of a distribution of the coin, this time it is interpreted only in a statistical expectation. This time can be ignored. The earlier discussion obviously applies here fully.

Despite the changed interpretation of the strategies. To restate briefly, it seems that many people faced with this matrix would choose a "bad" strategy, under which they might win or lose 10¢, rather than their "good" strategy, which would give them exactly 0¢. As Von Neumann and Morgenstern conclude some ambiguity in this phrase terms "good" and "bad." As they point out in connection with an equivalent matrix, suppose that a player should play a "bad" strategy: "If the opponent played the good strategy, then the player's mistake would not matter." (p. 134). Then in what sense would the "bad" strategy be a mistake? If A should pick H and B should pick H-H, would A's 10¢ win measure the "badness" of his "mistake"? True, A would have exposed himself to the loss of 10¢; the question is whether he should be called into account for doing so.

So much has been pointed out already. A new question that

suggests itself is: Why bother to play the game at all, if one prefers the certainty of 0 to the chance of winning or losing. The answer which Oskar Morgenstern once gave to this question was that in many situations one must play a game, even against one's wishes.

~~It may be that~~ The entire orientation of game-theory is implicit in this reply. If we should suppose that the game-models we have been studying represent uncertainty-situations in which an individual is forced, against his will, to make decisions, the authors' rationale for the minimax principle suddenly becomes much more convincing. The behavior of their "rational" player can indeed be described as that of a man whose sole concern is to come out of the game with as little loss as possible. This is not the attitude of a ^{actually} man/matching pennies, nor playing any game at all for entertainment or profit. It is, rather, the psychology of a timid man pressed into a duel.

~~When~~ When this point of view has been adopted, an admission by von Neumann and Morgenstern leaps to the eye:

"While our good strategies are perfect from the defensive point of view, they will (in general) not get the maximum out of the opponent's (possible) mistakes--i.e., they are not calculated for the offensive." (my italics. p. 164. Note use of the word "mistakes.")

This statement is absolutely decisive in determining the exact significance ~~and limitation~~ of the theory. Yet so casually is it introduced, so swiftly left behind, that it is not surprising that no published commentary has ~~mentioned~~ noted the passage.

After this one-sentence nod to the basic limitation of the theory, the authors immediately point out: "It should be remembered, however, that...a theory of the offensive, in this sense, is not possible without essentially new ideas." (p. 164). This may not be a recommendation of the old ideas. The authors have been distinctly

in failing to develop fully the implications of their concession. Is it not possible that what they term a "theory of the offensive" is precisely what would appeal to many readers as a theory of rationality? What is the justification of identifying "rational behavior" uniquely with "the defensive point of view"?

A passage by Fellner is very pertinent~~x~~ here:

"By doctoring the concept of profit maximization it would be possible to arrive at a theoretical construction in the framework of which a policy of maximum safety margins could be called a variant of the policy of profit maximization. We should merely have to define the expected profits (which are maximized) not as best-guess profits but as the profits which are expected in the even that certain comparatively unlikely possibilities materialize...If we use our concepts in this sense, then profit maximization becomes an unqualified axiom. But if this is done, some of the most essential problems of value theory are hidden so skillfully that it becomes difficult indeed to find them."

The redefinition of "expected profits" he describes is very close to the introduction of the principle of minimax in the game situation. It seems almost equally true in the latter situation that to interpret "maximizing security" as the unique form of "maximizing gain" under uncertainty is to obscure essential problems.

These conclusions are in marked contrast to the positive tone of the authors' remarks a few pages before they concede the defensive character of the theory:

"All this should make it amply clear that v' may ~~be~~ indeed be interpreted as the value of a play...There is nothing heuristic or uncertain about the entire argumentation...We have made no extra hypotheses about 'who has found out whose strategy' etc. Nor are our results for one player based upon any belief in the rational conduct of the other--a point the importance of which we have repeatedly stressed."). 160 the 'intelligence' of the players,

It would be hard not read into this that the "results" presented were of exceptionally general significance. Yet they actually rest on an ~~explicit~~ implicit assumption of a defensive psychology in the players, a temperament that is conservative almost to an unreasoning degree. Such players would pay no attention at all to the possibility of gains above those offered by the minimax strategy: which is to say, the

in failing to observe fully the implications of their conclusion.
It is not possible that what we have is "theory of the offensive"
is precisely what would appeal to many readers as a theory of re-
sistance. What is the justification of identifying "resistance" as
behavior "uniquely with" the defensive point of view?

This is of course

utter rubbish, when
applied to O_X of G
and mostly rubbish

when applied

elsewhere

great "existing strategy" as the unique form of "existing game"

understanding why it is so obvious essential processes.

These conclusions are in marked contrast to the positive tone
of the authors' remarks a few pages before they reached the de-

fective character of the strategy:

"All this should make it equally clear that we may be indeed be-

interpreted as the value of a play... There is nothing necessarily or
necessarily about the entire superstructure... We have said so much
hypothetically about 'who has found out whose strategy' etc. Now we

our results for one player based upon any belief in the rational be-
havior of the other--a point of indifference of which we have repeat-
edly stressed." (p. 160) the "intelligence" of the players.

It would be hard not to read this as the "results" presented were
of exceptionally general significance. Yet they actually rest on

an unrealistic assumption of a defensive hypothesis in its players,
a hypothesis that is counter-active almost to an understanding game.

Each player would say no other way at all to the possibility of gains
above those offered by the initial strategy: which is to say, no

rules which guide their behavior give no explicit attention to the possibility that other players might not follow the same rules.¹

xxxxxx It seems doubtful that the theory presented by vonNeumann and Morgenstern really does pass the criteria for a general solution of rational behavior set up by the authors themselves.

Percy Bridgman has made a comment on mathematical theorizing that sums up many of the conclusions of this paper:

"In mathematics...In many situations you find the solution and then ~~in~~ set yourself the problem of finding the problem this is the solution of....This is a well known method and has yielded many solutions, but obviously it is not a very good method of getting the solution of any specific problem."

~~VonNeumann and Morgenstern~~

In the course of modifying their concepts and approach so as to conform to particular mathematical "solutions," von Neumann and Morgenstern seem to me to have lost sight of their original problem. In my opinion, ^{proof of the} ~~the~~ existence of a saddlepoint in all game matrices admitting mixed strategies does not constitute a general solution of the specific problem of rational behavior in a two-person game; nor does their argument indicate that any general solution can be found, in the sense of a uniquely reasonable choice of strategies. Their "value" of the game is not an outcome that will actually be attained by all or most ~~reasonable~~ reasonable people in a single play; in many games and with many people it may not be attained even after a sequence of plays. Nor can a case be made convincing to all reasonable people that they should always behave so as to attain it. In terms of our basic test, many people "otherwise reasonable" would not reject all decisions inconsistent with the von Neumann-Morgenstern rules even after deliberation.

Their "solution" is not even empirically relevant to the behavior of all "conservative" players. One must ask: How conservative are they ? How much potential gain are they willing to forego in order to be

What is conservative

rules which guide their behavior give no explicit attention to the possibility that other players might not follow the same rules. It seems doubtful that the theory is accepted by von Neumann and Morgenstern really to be the criteria for a general solution of rational behavior set up by the authors themselves. Percy Bridgman has made a comment on mathematical theorizing that sums up many of the conclusions of this paper:

"In mathematics... In many situations you find the solution and then in retrospect the problem of finding the solution is the solution of... This is a well known method and has yielded many solutions, but obviously it is not a very good method of testing the solution of any scientific problem."

XXXXXXXXXXXXXXXXXXXX
In the course of modifying their concepts and approach so as to conform to particular mathematical "solutions," von Neumann and Morgen-

stein seem to me to have lost sight of the original problem. In my opinion, the existence of a solution in all games is a thing which does not constitute a general solution of the scientific problem of rational behavior in a broad sense. In their argument I indicate that any rational solution can be found, in the sense of a uniquely determined solution of a game. Their "value" of the game is not an outcome that will actually be attained by all or most persons who play in a single play; in many games not with any people if they not be attained with a sequence of plays. Nor can a case be made compelling to all reasonable people that they should always behave so as to attain it. In terms of our basic test, many people "otherwise reasonable" would not reject all solutions inconsistent with the von Neumann-Morgenstern rules even after reflection.

Their "solution" is not even explicitly relevant to the behavior of all "conservative" players. One must ask: how conservative are they? How much social capital are they willing to forego in order to be

You have argued quite skillfully that fools also exist / I agree with you

assured of a floor under their losses? Usually, the answer to that must be compared to the specific payoff function ~~matrix~~ Before any predictions are possible. Only if the answer is, "As conservative as it is possible to be" can the von Neumann-Morgenstern formula be used to predict the outcome without reference to the particular matrix

~~The problem that they have solved is~~

They have solved a very restricted problem. They develop fully in their preference-orderings the implications of the hypothesis that two players both (a) rank uncertain sets of outcomes according to the ~~mathematical~~ least **FALSE** element in each, and (b) rank probability distributions of outcomes according to the mathematical expectations. Given the particular conditions of information which they assume (i.e., knowledge by both of the permissible strategies and the payoff function) their hypotheses are empirically meaningful. They may well be also useful, both normatively and predictively: (1) a defensive or conservative policy is often desirable; (2) cautious pessimists do exist, whose behavior is consistent with the maximin principle in all situations. However, the restrictions on preferences stated in (a) and (b) seem overly special to be made general postulates of rationality.

Certainly the empirical, descriptive significance of the theory cannot be lessened if the assumptions are regarded as empirical hypotheses, whose relevance to particular situations is always to be tested. As it happens, there is reason to believe that the hypotheses will not always prove useful in describing the behavior of reasonable people in game-situations, nor is it always plausible ^{that players} ~~xxxxxxx~~ should be advised to conform against their inclinations. Whatever the remaining usefulness of the analysis, in the broad field of rational choice under uncertainty it cannot be said to provide a precise "theory...which gives complete answers to all questions." (p. 101)

a kind of a floor under their losses? Usually, the answer to that
 must be compared to the special report function which may
 operations are possible. Only if the answer is, "a conservative
 as it is possible to do" can the von Neumann-Morgenstern formula be
 used to predict the outcome without reference to the particular matrix

XXXXXXXXXXXXXXXXXXXX
 They have solved a very restricted problem. They have given fully
 in their preference-ordering
 the justification of the hypothesis that two players both (a) want

uncertain sets of outcomes according to the XXXXXXXXXXXX theory
 elements in each, and (b) rank probabilistic distributions of outcomes
 according to the rationalized expectations. Given the particular
 conditions of information which they assume (i.e., knowledge by both
 of the particular state, also and the payoff function) - all hypo-
 theses are empirically meaningful. They may well be also useful,
 both normatively and predictively; (1) a defensive or conservative
 policy is often desirable; (2) cautious pessimists do exist, whose
 behavior is consistent with the maximin principle in all situations.
 However, the restriction on preferences stated in (a) and (b) seem

4
 2
 5
 6
 8
 14
 16

Certainly the empirical, descriptive justification of the theory
 cannot be based on the assumptions are regarded as empirically
 correct, whose relevance to particular situations is always to be

3
 24
 will not always prove useful in describing the behavior of responsible
 that players
 people in game-theory, nor is it always plausible XXXXXXXXXXXX
 should be advised to conform against their inclinations. Whatever

the remaining usefulness of the analysis, in the broad field of
 rational choice under uncertainty it cannot be said to provide a
 precise
 "theory...which gives answers to all questions." (p. 101)

point out, it seems paradoxical to put the "danger of one's strategy being found ~~by~~ out by the opponent into an absolutely central position (p. 147) if the possibility of observation over a long series of plays is rejected. Their answer to this is that if a theory as determinate as the one they seek did exist, then the player would have to assume that his strategy had been found out, so the possibility of being found out would be present even under a (satisfactory) static theory. But the tenor of our discussion~~x~~ so far has been that no theory so determinate has been produced¹; so the paradox remains.

~~xx~~ Von Neumann and Morgenstern place great weight on the mathematical tradition sanctioning this approach, i.e., the logical derivation of properties of a solution on the assumption that a solution exists. But in the absence of an existence-theorem demonstrating the existence of a solution, the properties derived may be wholly useless or absurd. To illustrate the logical necessity of ascertaining the existence of a solution, Courant and Robbins cite the following fallacy: "1 is the largest integer. For let us denote the largest integer by x . If $x > 1$, then $x^2 > x$, hence x could not be the largest integer. Therefore x must be equal to 1." ("What is Mathematics?" Oxford University Press, 1941, p. 367.)

Von Neumann and Morgenstern certainly cannot show as a matter of logical necessity that an~~x~~ acceptable solution to the problem they pose exists, and the criticisms in this paper suggest that they have failed to produce a convincing case ~~for~~ for its existence ~~xxxx~~ on empirical or intuitive ~~xxxx~~ grounds. If we should decide that no satisfactory definition of a unique principle of rationality exists, then their arguments as to the necessary properties of such a principle become pointless.

Like argument about equil. properties, without discussion of stability.

There has been great confusion over the relation of mixed strategies to the minimax principle. The minimax principle could be advised, or used, whether or not the payoff function had a saddlepoint. ~~(Although xxxxx Neumann and Morgenstern might xxxxxx and~~
~~xxx~~ In fact, many, if not all, players who favored the minimax principle would probably be quite indifferent to the existence or not of a saddlepoint; hence the consideration of mixed strategies, which guarantee the existence of a bilinear form with a saddlepoint, is not necessary to the use of the minimax principle, whether or not